Abstracts

Tutorials

Tutorial 1 - *"Who Gets What? Fair Division of Indivisible Goods"*

Kurt Mehlhorn, Director Emeritus of the *Max Planck Institute for Informatics, The Max Planck Society for the Advancement of Science, Germany,* and Senior Professor of Computer Science, *Saarland University, Germany*

Abstract: A set of indivisible goods, e.g., a car, a house, a toothbrush, . . . , has to be split among a set of agents in a fair manner. Each agent has its own valuation function for sets of goods. What constitutes a fair allocation? When does a fair allocation exist? If it exists, can we compute it efficiently? Can we approximate fair allocations? There are three main notions of fairness in the literature: Envy-based, Share-based, and maximum Nash Social Welfare. I will discuss all three notions. In the first lecture, I will discuss what is known for the different notions of fairness. In the second lecture, I will discuss envy-free allocations in more detail. One of the notions of fairness is envy-freeness: nobody should get more than I do. For indivisible goods, envy-freeness cannot be achieved in general. Think of two person and one good which both persons like. The good has to be given to one of the persons and the other person will envy. Envy-freeness up to any good (EFX) is a relaxation. One person may envy another person, but upon the removal of any good from the other person's bundle, the envy goes away. EFX is a strong requirement as the following hypothetical dialogue shows. A person says to his brother. You got a house, a car, and a toothbrush. I envy you, but this is OK, because if I discard the toothbrush, I do not envy you anymore. Do EFX always exist?

Tutorial 2 - *"Total Search Problems in Algorithmic Game Theory"*

Paul Goldberg, Professor of Computer Science, *Oxford University, UK*

Abstract: Total search problems are problems of computing a solution that is known to exist but may be hard to find in the first place. This can happen when solutions have a "computationally inefficient proof of existence": an existence proof that does not come with an efficient algorithm. Various problems in game theory have this property, the best-known being the problem of computing a Nash equilibrium. Their complexity has been classified with complexity classes such as PPAD, PPA, PLS, and CLS. I will introduce these complexity classes and explain why they give evidence of computational hardness. I will give an overview of some of the results and proof techniques.

Tutorial 3 - *"Online Learning & Markets"*

Stefano Leonardi, Professor at the Department of Computer, Control & Management Engineering, *Sapienza Università di Roma, Italy* and **Federico Fusco**, Assistant Professor, *Sapienza Università di Roma, Italy*

Abstract: The last decades have witnessed a steep increase in the use of machine learning algorithms in countless economic applications, such as online advertising markets, financial exchanges, and energy markets. Learning algorithms are regularly deployed to operationalize decision-making procedures in scenarios characterized by (learning) agents that make strategic decisions by repeatedly interacting with complex environments. Motivated by these phenomena, a recent line of research at the intersection of online learning and economics investigates canonical economic models from the perspective of regret minimization. This tutorial introduces this research area, presenting the current results, the technical toolbox, and future directions. In particular, we investigate pricing, first and second-price auctions, and bilateral trade.

Talks

Talk 1 - *"Proportional Fairness in Non-Centroid Clustering"*

Ioannis Caragiannis, Professor at the Department of Computer Science, *Aarhus University, Denmark*, and Head of the *Computational Complexity & Game Theory research group*

Abstract: We revisit the recently developed framework of proportionally fair clustering, where the goal is to provide group fairness guarantees that become stronger for groups of data points (agents) that are large and cohesive. Prior work applies this framework to centroid clustering, where the loss of an agent is its distance to the centroid assigned to its cluster. We expand the framework to non-centroid clustering, where the loss of an agent is a function of the other agents in its cluster, by adapting two proportional fairness criteria -the core and its relaxation, fully justified representation (FJR)- to this setting. We show that the core can be approximated only under structured loss functions, and even then, the best approximation we are able to establish, using an adaptation of the GreedyCapture algorithm developed for centroid clustering [Chen et al., ICML 2019; Micha and Shah, ICALP 2020], is unappealing for a natural loss function. In contrast, we design a new (inefficient) algorithm, GreedyCohesiveClustering, which achieves the relaxation FJR exactly under arbitrary loss functions, and show that the efficient GreedyCapture algorithm achieves a constant approximation of FJR. We also design an efficient auditing algorithm, which estimates the FJR approximation of any given clustering solution up to a constant factor. Our experiments on real data suggest that traditional clustering algorithms are highly unfair, whereas GreedyCapture is considerably fairer and incurs only a modest loss in common clustering objectives.

Talk 2 - *"Utilitarian Distortion with Predictions"*

Alexandros Voudouris, Associate Professor at the School of Computer Science and Electronic Engineering, *University of Essex, UK*

Abstract: We study the utilitarian distortion of social choice mechanisms under the recently proposed learning-augmented framework where some (possibly unreliable) predicted information about the preferences of the agents is given as input. In particular, we consider two fundamental social choice problems: single-winner voting and one-sided matching. In these settings, the ordinal preferences of the agents over the alternatives (either candidates or items) is known, and some prediction about their underlying cardinal values is also provided. The goal is to leverage the prediction to achieve improved distortion guarantees when it is accurate, while simultaneously still achieving reasonable worst-case bounds when it is not. This leads to the notions of consistency and robustness, and the quest to achieve the best possible tradeoffs between the two. We show tight tradeoffs between the consistency and robustness of ordinal mechanisms for single-winner voting and one-sided matching, for different levels of information provided as prediction. *(Joint work with Aris Filos-Ratsikas and Georgios Kalantzis)*

Talk 3 - *"Online Combinatorial Allocation with Interdependent Values"*

Rebecca Reiffenhäuser, Assistant Professor for Theoretical Computer Science, *University of Amsterdam, The Netherlands*, and Member of the *Institute for Logic, Language and Computation*

Abstract: Combinatorial allocation problems and auctions are a main line of research in online algorithms and mechanism design, motivated by their natural and important applications. So far, they have been investigated largely under the assumption that buyers, who arrive online, are independent of one another. We instead tackle online combinatorial allocation problems under the interdependent values model, initially introduced by Milgrom (1982), where buyers' preferences depend on the opinions (or so-called signals) of others. We assume that buyers arrive in uniformly random order, also known as the secretary model. We show constant-competitive secretary algorithms and truthful mechanisms for the according, most general classes of buyer valuations where such are known without interdependence, given that valuations are additionally subadditive/XOS over the signals. Our results match the best guarantees known for the according single-choice problems with interdependence. *(Joint work with Michal Feldman, Simon Mauras, and Divyarthi Mohan (EC 2025))*

Talk 4 - *"Improved Last-Iterate Convergence Rates for Bilinear Zero-Sum Games"*

Michail Fasoulakis, Lecturer (Assistant Professor) in Theoretical Computer Science, *Royal Holloway, University of London, UK*, and Affiliated/Corresponding Scientist of Abroad at *ICS-FORTH, Greece*

Abstract: The recent years have seen a surge of interest in algorithms with last-iterate convergence for 2-player games, motivated in part by applications in machine learning. Driven by this, we revisit a variant of Multiplicative Weights Update (MWU), defined recently and denoted as Forward Looking Best Response MWU (FLBR-MWU). These dynamics are based on the approach of extra gradient methods, with the tweak of using a different learning rate in the intermediate step. So far, it has been proved that this algorithm attains asymptotic convergence but no explicit rate has been known. We answer this open question by establishing a geometric convergence rate for the duality gap. In particular, we first show such a rate, of the form O(c^t), till we reach an approximate Nash equilibrium, where c is independent of the game parameters (and c<1). We then prove that from that point onwards, the duality gap keeps getting decreased with a geometric rate, albeit with a dependence on the maximum eigenvalue of the Jacobian matrix.

Talk 5 - *"Airdrop Games"*

Paolo Penna, Research Fellow at *Input Output Global (IOG)*

Abstract: Launching a new blockchain system or application is frequently facilitated by a so called airdrop, where the system designer chooses a pre-existing set of potentially interested parties and allocates newly minted tokens to them with the expectation that they will participate in the system — such engagement, especially if it is of significant level, facilitates the system and raises its value and also the value of its newly minted token, hence benefiting the airdrop recipients. A number of challenging questions befuddle designers in this setting, such as how to choose the set of interested parties and how to allocate tokens to them.

Talk 6 - *"Algorithmic Monetary Policies for Blockchain Participation Games"*

Carmine Ventre, Professor of Computer Science, and Chair in Computational Finance, and Director of Informatics Finance Hub, and Interim Head of the Department of Informatics, *King's College London, UK*

Abstract: A central challenge in blockchain tokenomics is aligning short-term performance incentives with long-term decentralization goals. We propose a framework for algorithmic monetary policies that navigates this tradeoff in repeated participation games. Agents, characterized by type (capability) and stake, choose to participate or abstain at each round; the policy (probabilistically) selects high-type agents for task execution (maximizing throughput) while distributing rewards to sustain decentralization. We analyze equilibria under two agent behaviors: myopic (short-term utility

maximization) and foresighted (multi-round planning). For myopic agents, performance-centric policies risk centralization, but foresight enables stable decentralization with some volatility to the token value. We further introduce virtual stake -- a hybrid of type and stake -- to interpolate between proof-of-work and proof-of-stake selection. We prove that the initial virtual stake distribution critically impacts long-term outcomes, suggesting that policies must indirectly manage decentralization.

Talk 7 - *"Solving Neural Min-Max Games: The Role of Architecture, Initialization & Dynamics"*

Manolis Vlatakis, Assistant Professor in the Department of Computer Sciences, *University of Wisconsin-Madison, USA* & Affiliated Researcher at *Archimedes, Athena Research Center, Greece*

Abstract: Many core applications — like adversarial training and robust optimization — can be framed as zero-sum games between neural networks. Even though these games are non-convex and hard in theory, gradient methods often succeed in practice. In this work, we explain why: we show that overparameterization and hidden convexity can drive these dynamics to global equilibrium. Under mild conditions on initialization and network width, we prove convergence in two-layer neural net games — something not previously known.

Talk 8 - *"Learning with Systematic Bias and Imperfect Data"*

Alkis Kalavasis, Postdoctoral Fellow at the Institute for Foundations of Data Science (FDS), *Yale University, USA*

Abstract: In many applications in Econometrics and Statistics, we only have access to imperfect data due to systematic bias in the data collection process and the incentives of strategic agents. In this talk, we will present a general formulation for learning in the presence of such biasing mechanisms. Under this framework, we will discuss learning algorithms for inference under self-selection bias and coarsening, and estimation of treatment effects in observational studies.

Talk 9 - *"Solving Hidden Monotone Variational Inequalities with Surrogate Losses"*

Ioannis Mitliagkas, Associate Professor in the Department of Computer Science and Operations Research (DIRO), *Université de Montréal, Canada* and Staff Research Scientist at *Google DeepMind Montréal, Canada* and Affiliated Researcher at *Archimedes, Athena Research Center, Greece*

Abstract: Deep learning has proven to be effective in a wide variety of loss minimization problems. However, many applications of interest, like minimizing projected Bellman error and min-max optimization, cannot be modelled as minimizing a scalar loss function but instead correspond to solving a variational inequality (VI) problem. This difference in setting has caused many practical challenges as naive gradient-based approaches from supervised learning tend to diverge and cycle in the VI case. In this work, we propose a principled surrogate-based approach compatible with deep learning to solve VIs. We show that our surrogate-based approach has three main benefits: (1) under assumptions that are realistic in practice (when hidden monotone structure is present, interpolation, and sufficient optimization of the surrogates), it guarantees convergence, (2) it provides a unifying perspective of existing methods, and (3) is amenable to existing deep learning optimizers like ADAM. Experimentally, we demonstrate our surrogate-based approach is effective in min-max optimization and minimizing projected Bellman error. Furthermore, in the deep reinforcement learning case, we propose a novel variant of TD(0) which is more compute and sample efficient.

Talk 10 - *"On Recent Advances in Computational Complexity of Team Games and Beyond"*

Ioannis Panageas, Assistant Professor at the School of Information & Computer Science, *University of California, Irvine, USA* and Director of *The Games, Optimization, Algorithms, and Learning Lab (GOALLab)* and Lead Researcher at *Archimedes, Athena Research Center, Greece*

Abstract: We show that computing ϵ-Nash equilibria in 3-player adversarial team games -wherein a team of 2 players competes against a single adversary- is CLS-complete, resolving the complexity of Nash equilibria in such settings. Our proof proceeds by reducing from symmetric ϵ-Nash equilibria in symmetric, identical-payoff, two-player games, by suitably leveraging the adversarial player so as to enforce symmetry, without disturbing the structure of the game. In particular, the class of instances we construct comprises solely polymatrix games, thereby also settling a question left open by Hollender, Maystre, and Nagarajan (2024). We also provide some further results concerning equilibrium computation in adversarial team games. Moreover, we establish that computing symmetric (first-order) equilibria in symmetric min-max optimization is PPAD-complete, even for quadratic functions. Building on this reduction, we further show that computing symmetric ϵ-Nash equilibria in symmetric, 6-player (3 vs. 3) team zero-sum games is also PPAD-complete, even for $\epsilon = poly(1/n)$. As an immediate corollary, this precludes the existence of symmetric dynamics -which includes many of the algorithms considered in the literature- converging to stationary points. Finally, we prove that computing a non-symmetric $poly(1/n)$-equilibrium in symmetric min-max optimization is FNP-hard.

Talk 11 - *"Algorithmic Foundations for Contrastive Embeddings"*

Vaggos Chatziafratis, Assistant Professor of Computer Science & Engineering, University of California, Santa Cruz, USA and Affiliated Researcher at *Archimedes, Athena Research Center, Greece*

Abstract: The goal of contrastive embeddings (also known as ordinal embeddings) is to find convenient data representations that preserve the relative order between distances, and are well-motivated by applications in nearest-neighbor search, recommendation, ranking, and more recently,

by successful contrastive learning paradigms in training neural networks. The input typically consists of "anchor,positive,negative" triplets for 3 items (a,p,n) indicating that distance $||a-p|| < ||a-n||$ in the final embedding. Unfortunately, basic questions regarding dimensionality, approximation and/or convergence for contrastive embeddings are not well-understood. In this talk, we will give an overview of some recent algorithmic progress and also discuss connections to Phylogenetic Constraint Satisfaction Problems, such as Triplet Reconstruction methods in computational biology. No prior background is assumed, and the talk will be self-contained.

Organizing & Scientific Committee

- **Georgios Amanatidis**, Assistant Professor at the Department of Informatics of the *Athens University of Economics and Business (AUEB), Greece*, and Lead Researcher at *Archimedes, Athena Research Center, Greece*
- **Georgios Christodoulou**, Associate Professor in the School of Informatics at the *Aristotle University of Thessaloniki (AUTh), Greece*, and Lead Researcher at *Archimedes, Athena Research Center, Greece*
- **Dimitris Fotakis**, Professor at the School of Electrical and Computer Engineering, *National Technical University of Athens (NTUA), Greece*, and Lead Researcher at *Archimedes, Athena Research Center, Greece*
- **Evangelos Markakis**, Professor at the Department of Informatics of the *Athens University of Economics and Business (AUEB), Greece*, and Lead Researcher at *Archimedes, Athena Research Center, Greece*
- **Alkmini Sgouritsa**, Assistant Professor at the Department of Informatics of the *Athens University of Economics and Business (AUEB), Greece*, and Lead Researcher at *Archimedes, Athena Research Center, Greece*

Co-Organizers

ΟΙΚΟΝΟΜΙΚΟ ΠΑΝΕΠΙΣΤΗΜΙΟ ΑΘΗΝΩΝ — ATHENS UNIVERSITY OF ECONOMICS AND BUSINESS

ARCHIMEDES

ATHENA Research & Innovation Information Technologies